arXiv:1604.02171v1 [cs.OS] 7 Apr 2016

Aware: Controlling App Access to I/O Devices on Mobile Platforms

Giuseppe Petracca* *gxp18@cse.psu.edu*
Ahmad Atamli* *atamli@cs.ox.ac.uk*
Yuqiong Sun* *yus138@cse.psu.edu*
Jens Grossklags* *jensg@ist.psu.edu*
Trent Jaeger* *tjaeger@cse.psu.edu*
*The Pennsylvania State University
*University of Oxford

Abstract *Smartphones' cameras, microphones, and device displays enable users to capture and view memorable moments of their lives. However, adversaries can trick users into authorizing malicious apps that exploit weaknesses in current mobile platforms to misuse such on-board I/O devices to stealthily capture photos, videos, and screen content without the users' consent. Contemporary mobile operating systems fail to prevent such misuse of I/O devices by authorized apps due to lack of binding between users' interactions and accesses to I/O devices performed by these apps. In this paper, we propose Aware, a security framework for authorizing app requests to perform operations using I/O devices, which binds app requests with user intentions to make all uses of certain I/O devices explicit. We evaluate our defense mechanisms through laboratory-based experimentation and a user study, involving 74 human subjects, whose ability to identify undesired operations targeting I/O devices increased significantly. Without Aware, only 18% of the participants were able to identify attacks from tested RAT apps. Aware systematically blocks all the attacks in absence of user consent and supports users in identifying 82% of social-engineering attacks tested to hijack approved requests, including some more sophisticated forms of social engineering not yet present in available RATs. Aware introduces only 4.79% maximum performance overhead over operations targeting I/O devices. Aware shows that a combination of system defenses and user interface can significantly strengthen defenses for controlling the use of on-board I/O devices.*

1 Introduction

Nowadays, mobile platforms are equipped with cameras, microphones and wide screens, which enable a variety of popular functions, ranging from audio/video recording to displaying information to users. Many apps now utilize these functions to provide services that leverage on-board Input/Output (I/O) devices. For example, many apps now support voice and video messages, as well as photo/video shooting and editing. Even insurance or banking apps now utilize mobile platforms' cameras to collect sensitive information to expedite claim processing or check depositing [15, 17]. Apps able to record the screen content are also available for remote screen sharing or tutorial editing.

However, uncontrolled access to on-board I/O devices can enable malicious apps, with access to these devices, to exfiltrate sensitive information. Adversaries have built malware apps, referred to as *Remote Access Trojans* (RATs), that abuse authorized access to such devices to extract audio, video, screen content, and more, from mobile platforms, such as smartphones and tablets. Malware in the wild both surreptitiously records data from a variety of mobile devices [3, 1, 13], but also performs directed attacks, such as constructing three-dimensional models of indoor environments [22] and extracting credit card numbers and PIN numbers [23] from screenshots or keyboards' tones. Furthermore, uncontrolled accesses to on-board cameras and microphone can become sensitive if apps can stealthily take photos or videos and record audio by running a service in background[1]. Additionally, RAT apps can also use social-engineering techniques [16] to hijack user-requested activities, such as showing a soft-button on screen that supposedly allows the user to take a picture when in reality the user action will trigger voice recording instead through the smartphone's microphone.

Current defenses do not prevent malicious apps that happen to be granted access to I/O devices from exfiltrating sensitive data. Android, iOS, and Windows Phone OS all require users to authorize apps for access to I/O devices, such as the camera and microphone, at install time or at first use. In many cases, users may assume that the use of such devices will be important, if not fundamental, for the effective operation of such apps. Should the user authorize an app, the app can

[1] 75% of operations requiring permissions are performed when the screen is off or apps are running in background as services [36]

then choose when to use the device. In addition, access to screen content is not even mediated by any of the mobile operating systems. More restrictive security models, such as SE Android [29], cannot control apps access to such devices further, as they mainly protect the lower level Android system. Some research systems aim to prevent unauthorized access to resources, such as these devices, by empowering apps to assist in the decision making [25, 30, 31]. However, since apps may be malicious, attacks cannot be prevented by this method. Alternatively, researchers have explored auditing the use of such devices [34] and providing visible indication of app behaviors [28], but the former only detects attacks retroactively after the data has been exfiltrated, whereas user notification alone requires users to pay attention to each operation's security status, continuously, to avoid missing attacks.

In this work, we propose the Aware framework for authorizing app requests to perform operations using I/O devices, which binds app requests with user intentions to make all uses of I/O devices explicit. To do this, we take the following steps. First, app requests for operations using I/O devices are intercepted by the Aware-enabled services to mediate all security-sensitive operations. Second, Aware makes each app request visible to the user, independently from the app, to enable users to express their intents without being spoofed. Third, Aware makes ongoing operations, targeting I/O devices, visible to users, so they may choose to terminate the operation. As opposed to previous solutions, Aware does not depend on apps to govern users, but rather Aware links app requests and user input. Also, Aware maintains the security status of ongoing operations, so user actions are only necessary at operation initiation and completion, rather than requiring ongoing user monitoring [28].

We have implemented Aware on Android OS (android-6.0.1_r5 version) and found it to perform effectively, by adding a maximum overhead of 4.79% (minimum 2.19%). We have performed a user study, involving 74 human subject. Without Aware, only 18% of the study participants were able to identify attacks from tested RAT apps. Aware systematically blocks all the attacks in absence of user consent and enabled study participants to identify 82% of social-engineering attacks tested to hijack approved requests, including some more sophisticated forms of social engineering not yet present in available RATs. This paper makes the following contributions:

- We reverse-engineer two real-world RAT apps and two proof-of-concept RAT apps to systematically study and categorize the different techniques used by adversaries in mounting attacks targeting on-board I/O devices. We identify five security properties that must be satisfied in order to ensure protection against stealthy operations from malicious apps targeting on-board I/O devices.
- We propose Aware, a security framework, that introduces defense mechanisms for enforcing these five security properties by mediating app requests to I/O devices and matching those requests to user consent, which blocks all unapproved requests and maintains the security status of ongoing requests to the user to enable prevention of social-engineering attacks.
- We conduct an extensive user study involving 74 human subjects to evaluate: (1) users' awareness of RAT attacks targeting I/O devices, (2) effectiveness of RAT attacks targeting on-board I/O devices, and (3) effectiveness of our proposed defense mechanisms in increasing users' awareness and control over sensitive operations targeting I/O devices.
- We evaluate our approach on five RAT apps and eight widely-used apps, to show that it is possible to prevent against attacks from RAT apps without compromising functionality or introducing significant performance overhead.

2 Problem Definition

In this section, we describe Remote Access Trojans (RATs) to demonstrate attacks that exploit use of I/O devices. We then examine the state-of-the-art in permission granting for mobile platforms to understand why such malware apps are capable of exploiting I/O devices on smartphones. We then outline the challenges for defenses capable of blocking such attacks.

2.1 Remote Access Trojans (RATs)

On smartphones, *Remote Access Trojans* (RATs) are malicious apps users may be tricked into installing on their smartphones that aim to violate users' privacy and data confidentiality. RATs collect security-sensitive information through on-board I/O devices, such as cameras, microphones, and screen buffers using authorized app permissions to perform stealthy, malicious operations including taking photos and videos, recording audio, or capturing screenshots.

Researchers have designed and developed mobile RATs to demonstrate limitations of current access control models adopted in mobile OSs. Examples include: *PlaceRaider* [22], which uses the camera and other sensors to construct rich, three-dimensional, models of indoor environments; and *Soundcomber* [23], which exfiltrates sensitive data, such as credit card and PIN numbers, from both tones and speech-based interaction with phone menu systems. Real-world RATs are also available online for purchase. Two popular ones, widely discussed in security blogs and anti-virus companies, are:

Dendroid [1], which takes photos using the phone's camera, records audio and video, downloads existing photos, records calls, sends texts, etc. and *Krysanec* [3], which takes photos, records audio through the microphone, locates victims via devices' GPS, views victims' contact list, and intercepts and sends text messages.

We reverse engineered and statically analyzed these RATs. The two real-world RATs were leaked online, whereas the two proof-of-concept RATs have been shared by researchers. From the analysis, we obtained details on how RATs work. We found out that: (1) *all* analyzed RATs require a specific set of permissions to access on-board I/O devices, except when accessing the screen buffers; (2) *all* of them run a background service that stealthily performs malicious operations to abuse permissions granted at install time or first use; (3) *none* of their activities is shown on screen; (4) *all* of them need access to the Internet to leak security-sensitive data collected; and (5) their stealthy operations are *never* associated with any user interaction with the smartphone.

2.2 Limits of Permission Granting

Mobile OSes currently support two default mechanisms to grant apps permission to access on-board I/O devices.

First, in Android OS, users grant apps permission to access I/O devices at *install time*. Apps receiving permission at install time can then access I/O devices *at any time*, without further user approval, so users are unable to track *how* and *when* sensitive on-board I/O devices are accessed by apps at runtime. Table 1 summarizes the permission sets required by the analyzed RAT apps to perform stealthy operations. We found out that: (1) *all* permissions used by RAT apps to perform stealthy operations are classified as *dangerous* by the official Android OS documentation [4]; (2) users are *never* notified about accesses to security-sensitive on-board I/O devices[2]at runtime, via on-screen prompts, or notifications; and (3) the *same* sets of permissions are used by well-known benign apps downloaded by millions of users worldwide. We performed an extensive analysis of permission sets used by apps by randomly selecting 74 apps from third-party app stores [11, 10] and 329 apps from the official Google Play [8]. The results cause concern: 83.89% of apps from the official Google Play store could potentially take stealthy screenshots, whereas 25.68% of apps from third-party app stores could potentially take stealthy photos (complete analysis results are reported in Appendix A).

Second, starting from Android OS 6.0 (Marshmallow) and in other mobile operating systems, such as Apple iOS and Windows Phone OS, users are prompted

Table 1: Android Permissions used by RAT apps

	Permissions required to perform Operation	Protection Levels	User Notified
Stealthy Photo	CAMERA WRITE_EXTERNAL_STORAGE	Dangerous	No
Stealthy Video	RECORD_AUDIO, CAMERA WRITE_EXTERNAL_ STORAGE	Dangerous	No
Stealthy Audio	RECORD_AUDIO WRITE_EXTERNAL_STORAGE	Dangerous	No
Stealthy Screenshot	WRITE_EXTERNAL_STORAGE	Dangerous	No
Remote Data Transfer	INTERNET	Dangerous	No[1]

with access requests *the first time* apps request access to I/O devices. Researchers have shown that, while these prompts attempt to verify users' intent, in practice, they create an excessive burden on users, which leads to users ignoring these prompts eventually[3]. In addition, these mobile operating systems allow users to manage permission grants at runtime by accessing a per app or per I/O device permission control panel. This feature allows for better flexibility in permission granting, since it is now possible to revoke, at runtime, permissions granted to apps at install time or first use.

Unfortunately, neither of these mechanisms ensure that sensitive operations targeting on-board I/O devices are performed *only* in response to users' interaction with running apps, which *must unmistakeably bind the user's consent with specific security-sensitive operations targeting an I/O devices*. In the absence of such binding, malicious apps are free to leverage I/O devices, even while running as background services, once they can trick users into granting them permissions as shown in Figure 1.

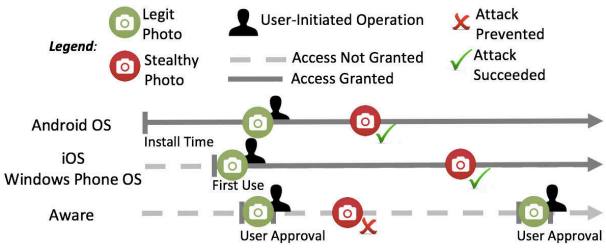

Figure 1: Effect of the lack of binding between users' interaction and operations performed by running apps

2.3 Challenges in Preventing RAT Exploits

Researchers have identified several attacks targeting on-board I/O devices [22, 23], and proposed various defense mechanisms [25, 26, 29, 5]. Unfortunately, significant attacks are still capable of circumventing proposed defenses. Furthermore, most anti-malware tools available for smartphones are not able to identify apps behaving as RATs, especially if not advertised and commercialized as

[2]Users could notice the Wi-Fi or cellular network icon on the phone screen status bar, but they do not know what app is responsible for the network traffic, and what data is flowing out through the network.

[3]Prompts are disruptive and cause excessive fatigue, conditioning user to simply accept any prompt query, resulting in undermining the usefulness of the prompts [24, 19, 36].

spying tools on the Web[4]. Similarly, current static analysis tools [12, 32, 33, 14] designed to analyze apps' source code to identify malice are often not able to find stealthy operations targeting I/O devices[5].

On one hand, several attempts have been made, in the past, to include users or surrounding environments in the access control decision mechanism. Unfortunately, User-Driver Access Control (UDAC) [25] is subject to social-engineering attacks [28], whereas World-Driven Access Control (WDAC) [26] limits the objects that may be recorded, but does not enable users to control when apps use I/O devices for recording. However, solutions to prevent social-engineering attacks [28] currently require users to pay attention to the security status of an operation throughout its execution, rather than just at the beginning and end. On the other hand, researchers have also proposed methods to augment contemporary access control models based on the Android Permission mechanism [4, 24] and SE Android [29, 5]. Android Security Modules (ASM) [30] enable apps to assist in security decision-making, but unfortunately the apps performing operations are the adversary we must control; whereas Android Security Framework (ASF) [31], which operates at the kernel level, does not have the necessary information about higher level events required to detect security-sensitive operations performed by processes running apps. Additionally, solutions that leverage hardware support for app isolation, such as Samsung Knox [18], would prevent apps in one domain from stealing sensitive data from apps in other domains, but RAT apps can still perform stealthy operations within their own domain.

Despite the various efforts above, several open challenges remain to be addressed:

- Once apps obtain permission, at *install time* or at *first use*, they may stealthily access sensitive I/O devices *at any time*.
- Mobile operating systems lack a method to *connect user interactions with security-sensitive operations* targeting on-board I/O devices for controlling access to such operations.
- Apps may use *social-engineering* techniques [16, 27, 28] to hijacking user-intended activities and trick users in authorizing undesired operations.

In addition, any effectitve solution to these problems must only require user input and attention consistent with normal application use.

[4]We have tested the 15 most popular Android anti-malware tools, complete results are reported in Appendix B.

[5]We have tested 2 static and 2 dynamic analysis tools currently adopted by researchers and the general mobile app community. Complete results are reported in Appendix C.

3 Background

In this paper, we focus our attention on Android OS due to its open-source nature, availability, and popularity [2]. Similar considerations hold true for other mobile operating systems, such as Apple iOS and Windows Phone OS. In the following subsections, we provide background information useful to understand the mechanisms proposed in this paper.

3.1 I/O Device Management in Android

We briefly describe how processes obtain access to on-board I/O devices in Android OS. For performance and security reasons, only the Linux kernel has direct access to on-board I/O device drivers. The Hardware Abstraction Layer (HAL) implements an interface that allows system services (privileged Android processes) to indirectly access on-board I/O devices via well-defined APIs. SE Linux for Android guarantees that only system services can access on-board I/O devices at runtime. Thus, apps must communicate with system services, through the Binder mechanism [4], to request execution of specific operations targeting on-board I/O devices. The system services designed to handle requests from apps would then execute the operations on behalf of processes running apps, if and only if, the necessary permissions have been granted to the requesting apps by the user or operating system. Permissions are validated by the *Package Manager*, part of Android OS, each time apps request to perform operations targeting I/O devices. For instance, for operations targeting the microphone, the *Package Manager* checks the apps `AndroidManifest.xml` files to verify if the RECORD_AUDIO permission has been granted to the requesting app. If the required permissions are not granted to apps, the *Package Manager* fires up a security exception to communicate the operation abortion. Android services do not require a permission check for apps to access the screen buffer.

3.2 Android UI Graphical Elements

The Android User Interface is composed by three main graphical elements, as depicted in Figure 2 (A). Different manufacturers may different layouts, but all distributions of the Android OS have the same three main elements.

The *Status Bar* shows the device's state, such as battery level and network connectivity. The *Navigation Bar* includes three navigation buttons to interact with all currently running apps and the home screen. The *Activity Window* is the only portion of the screen that processes running apps can draw graphical elements on, such as *Activities* and *Views* inside activities. All activities created by apps are drawn in the system-managed *Activity Window*. Activities are organized in a stack managed by the *ActivityManager* system service, the only process al-

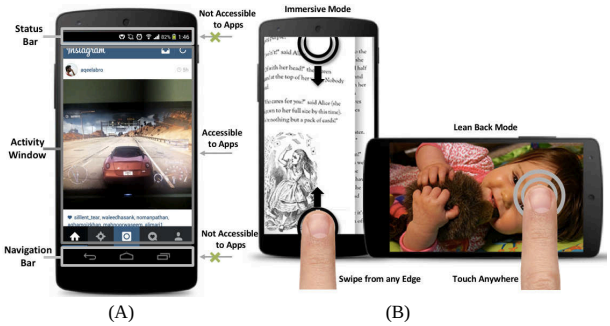

Figure 2: Android User Interface Graphical Elements and Gestures to bring back the Navigation and Status Bars into view when apps are in full screen mode

lowed to manage these three main graphical elements. The only activity shown to the user is the one on top of the stack, even though, previously displayed activities might be visible if the new activity on top is only partially covering the *Activity Window*.

The *Activity Window* becomes the only graphical element visible on screen when apps go in *Full Screen* mode. Starting from version 4.4, the Android OS offers apps two approaches to go full screen: *Lean Back*[6] and *Immersive*[7]. In both approaches, all persistent system bars are hidden. The difference between them is how the user brings the bars back into view, as shown in Figure 2 (B). In *Lean Back* mode, users can bring back the bars by simply touching the screen anywhere. Whereas, in *Immersive* mode, they can just swipe from any edge where a system bar is hidden. This gestures are easy and intuitive, an explanatory message appears on screen the first time the app goes full screen.

4 Security Model

When designing Aware, we focus on defending against any process running an app that has permissions to access security-sensitive, on-board I/O devices (e.g., camera, microphone, and touch screen) based on the following threat model.

We assume that adversaries have no control over the operating system (e.g., Linux kernel and Android OS) or Android system apps (e.g., SystemUI) and services (e.g., Binder, AudioSystem, MediaServer, SensorManager, InputManager and WindowManager). Next, we assume that SEAndroid enforces mandatory access control over access to on-board I/O devices, preventing unauthorized access from apps containing native code. Therefore, we assume that only system services can access on-board I/O devices indirectly through the use of the Java Native Interface (JNI) [7] to the Linux kernel. Further, we assume apps can access I/O devices *only* through APIs provided by the standard Android SDK [4]. Finally, we assume that Android system services can enforce access control policies that are applied at the time that data would be collected from an I/O device only. We aim to control whether apps can receive data produced by on-board I/O devices, but do not provide any guarantee after the app is granted access to the data itself.

Adversaries in control of apps may cause threats by executing the following operations. First, adversarial apps request use of security-sensitive operations on on-board I/O devices, including accessing devices' camera to take picture or video recording, devices' microphones to record users' voices or surrounding environments, and devices' screens to steal security sensitive information displayed to users. Additionally, adversarial apps may alter the display to launch social-engineering attacks to trick users into consenting to operations they do not want by overlaying a frame buffer over another apps' or system component's display or displaying a request for one operation then performing a different operation.

5 Aware Design Overview

5.1 Aware Operation

An overview of the Aware design is shown in Figure 3. In the overview description, we use Android OS as reference, similar considerations hold for other mobile operating systems available on the market. Typically, a process *Prc*, running an Android app, sends a request to perform an operation *Opr* using a specific I/O device *Dev* (step ❶). An example could be a request from an app (*Prc*) to access the front camera (*Dev*) and take a photo (*Opr*). The request is received by a system service *Srv* (e.g., the MediaServer for the camera), one of the privileged processes in charge of authorizing and processing access requests to system resources. At first, the conventional access control mechanisms are activated (step ❷). For instance, in Android OS, the Android permissions and the SE Android access control policy checks are activated. If the result of the conventional access control enforcement is a denial, then the system service *Srv* is notified of the security exception, which in turn notifies the requesting process *Prc* of the access denial (step ❸). Otherwise, if the result of the conventional access control enforcement is a grant, Aware is activated in order to implement its additional access control conditional checks.

The Aware Conditional Engine collects information about the process *Prc* requesting a certain operation *Opr* over a target device *Dev* (step ❹) through the system service *Srv*. Based on this information, the Aware Conditional Engine then identifies and selects, from the Conditional Rule Store, the set of conditional rules *Cnds* that must be satisfied to allow operation *Opr* to be performed

[6] Used when users won't be interacting heavily with the screen while consuming content, like while watching a video.

[7] Mainly intended for apps in which the user will be heavily interacting with the full screen as part of the primary experience, like while playing games, viewing images in a gallery, or reading a book.

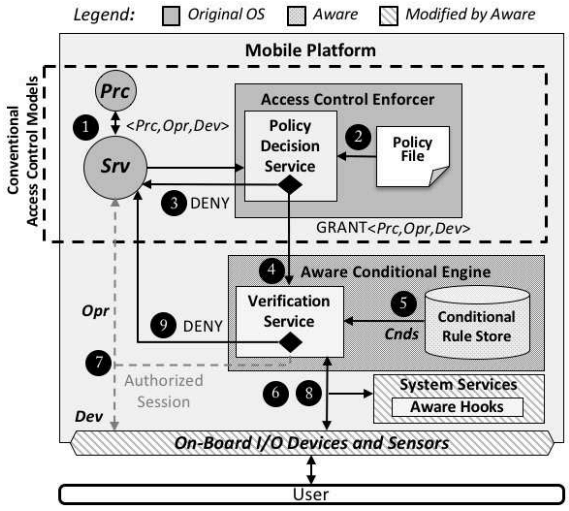

Figure 3: Overview of the Aware Design

on behalf of process *Prc*, over the target device *Dev*, by the system service *Srv* (step ❺). Subsequently, Aware collects the measurements necessary to evaluate if the selected conditional rules are satisfied (step ❻). Measurements can involve readings from on-board I/O devices and sensors, or system events necessary to verify specific environmental conditions. If and only if *all* the selected conditional rules *Cnds* are satisfied, then Aware generates an authorized session during which the system service *Srv* is authorized to perform the security-sensitive operation *Opr*, over the target device *Dev*, on behalf of process *Prc* (step ❼). During authorized sessions, Aware verifies that all conditional rules remain satisfied during the entire session and notifies users about ongoing operations (step ❽). Whenever *any* of the selected conditional rules is not satisfied, Aware abort the operation *Opr* and notifies the service *Srv* about the set of unsatisfied conditional rules via a denial message (step ❾).

5.2 Aware Conditional Engine

Aware Conditional Engine uses *conditional rules* to authorize app requests to I/O devices and maintain security state during any authorized sessions. To verify the satisfaction of conditional rules, Aware is designed to collect inputs from I/O devices, sensors and system services at run-time, by using additional hooks placed inside the Android framework. Aware conditional rules are of three types: (1) *Preconditions* that must be satisfied before an operation targeting I/O devices is authorized; (2) *Ongoing Conditions* that must be satisfied while the authorized operations targeting I/O devices are being performed, until their completion; and (3) *Exit Conditions* that must be all satisfied once authorized operations terminate due to users' actions or runtime exceptions.

Aware *Preconditions* ensure two security properties. **SP1** All app requests to perform security-sensitive operations targeting on-board I/O devices must be authorized. This security property prevents processes from performing such operations stealthly. **SP2** Security-sensitive operations performed using on-board I/O devices match a users' consenting action. This property ensures that every such operation is initiated by a user action. Aware *Ongoing Conditions* ensure two additional security properties. **SP3** Ongoing security-sensitive operations targeting I/O devices are always visible to users. This security property enables users to check that the authorized operation is what they expected, reducing the possibility of undetected social-engineering attacks. **SP4** Ongoing security-sensitive operations targeting I/O devices are always logged. Such logging enables users to examine the progress of ongoing security-sensitive operations, which may enable termination of an ongoing operation deemed suspicious or retrospective analysis of past operations. Aware *Exit Conditions* ensure one more security property. **SP5** All ongoing security-sensitive operations targeting I/O devices are visible to the user as long as they run. This property ensures that apps cannot keep operations running after user terminates them and the user interface correctly removes only terminated operations from the display. The conditional rules and security properties above mentioned are summarized, for future reference, in Table 2. To express Aware conditional rules, the Usage Control (UCON) model [20] could be adopted.

6 Aware Design

We present the Aware design in terms of four mechanisms necessary to fulfill the five security properties.

6.1 Mediation of Access Requests

Complete mediation of *all* requests to access I/O devices from processes running apps and matching each request to a user input corresponding to the app request is necessary to ensure that only authorized app requests are run, guaranteeing **SP1**. Mediation involves several system services, such as services controlling gestures on screen, or handling requests from running processes, when targeting on-board I/O devices and sensors. Users' interaction events have to be aggregated and mapped to requests instantiated by processes running apps. The SE Android reference monitor must be extended to ensure *complete mediation* of all security-sensitive operations targeting I/O devices [21]. Identifying the right locations where to place additional hooks, to mediate every access to I/O devices, is challenging.

Mainly, complete mediation of accesses to I/O devices can be achieved in two ways: (1) by placing hooks inside the Android framework and libraries, or (2) by placing hooks inside the Linux kernel and I/O device drivers. To achieve complete mediation of accesses to I/O devices, which are low-level system resources, the kernel and de-

Table 2: Aware Conditional Rules and Security Properties

Preconditions	Security Properties
P1 User interacts with process *Prc* to request operation *Opr* targeting device *Dev* P2 Process *Prc* requests Service *Srv* to perform operation *Opr* over device *Dev* P3 User is aware of what operation *Opr*, targeting device *Dev*, is going to be performed by Service *Srv* on behalf of process *Prc* P4 User approves operation *Opr* on behalf of process *Prc* targeting device *Dev*	SP1 All app requests to perform security-sensitive operations targeting on-board I/O devices must be authorized SP2 Security-sensitive operations performed using on-board I/O devices match a users' consenting action
Ongoing Conditions	
O1 User has continuous visibility of operation *Opr* performed on behalf of process *Prc* over device *Dev* O2 The authorized session *Ses*, for process *Prc* to execute operation *Opr* over device *Dev*, is logged to allow retroactive actions	SP3 Ongoing security-sensitive operations targeting I/O devices are always visible to users SP4 Ongoing security-sensitive operations targeting I/O devices are always logged
Exit Conditions	
E1 The authorized session *Ses*, relative to process *Prc*, terminates E2 Termination of session *Ses* is logged E3 User has visibility that session *Ses* has been terminated	SP5 All on-going security-sensitive operations targeting I/O devices are visible to the user as long as they run

vice drivers seem to be the most appropriate place where to add mediation hooks. However, two main issues arise from this approach: (1) low level hooks would not have the required level of information to map requests to processes running apps, due to the fact that requests are always handled by system services on behalf of the requesting processes; (2) mobile platforms are equipped with different I/O devices, which would require the operating system to be able to support customized hooks defined for different drivers by driver vendors.

In Aware, hooks are placed at the Android framework and libraries level, to avoid the above mentioned issues. Aware *Hooks* provide complete mediation, because system services are the only path through which processes, running apps, can access I/O devices and sensors, due to Android framework architecture and MAC rules enforced by SE Android [5]. We have dynamically analyzed the Android framework and libraries code, relative to SDK APIs handling accesses to I/O devices and sensors, to validate complete mediation, and check that every access to the I/O devices and sensors is captured by one of the 18 hooks introduced. Retaining such logging could be used to detect errors, if any exist. Callbacks from hooks inform the Aware Conditional Engine about users interacting with processes running apps, and requesting operations over I/O devices. These callbacks are used to validate precondition P1. Aware *Hooks* also capture resources acquisition and release by system services operating on behalf of processes running apps. Callbacks from these hooks are used to validate precondition P2 and exit condition E1. Satisfying these conditions is sufficient to reliably bind users interaction with apps requests to operate over I/O devices, therefore guaranteeing SP1.

6.2 Visibility over Sensitive Operations

Secure display of operations targeting I/O devices when they are requested, when they are ongoing, and when they are terminated is necessary to fulfill multiple guarantees. First, by maintaining visibility of operations after they are authorized, users may identify undesired operations approved by mistake to guarantee SP3. Furthermore, ensuring that operations are visible to users as long as they run guarantees that there are no stealthy operation on I/O devices ongoing SP5. Visibility over accesses to I/O devices from running apps may be provided to users in four different ways: (1) via notification lights, similar to those used for cameras on laptops or external USB cameras; (2) by playing a distinctive sound, similar to the shutter sound produced when taking a photo; (3) by displaying notification icons, similar to the location icon shown on the status bar; and (4) by visualizing alert messages on screen. Unfortunately, notification lights, sounds, or notification icons can only alert users about accesses to sensitive I/O devices, but cannot convey exact information about operations performed, target devices, and processes responsible for such operations. Furthermore, the sounds might not be audible in silent or vibrate mode. A better way to convey complete information about operations performed over I/O devices, by running processes, is by displaying on screen alert messages to users.

Solutions that make use of the *Activity Window* portion of the screen, to display access notifications or alert messages, are subject to user deception attacks, were screen overlays are used by malicious apps to surreptitiously replace, or mimic, the GUI of other apps and mount social-engineering attacks, or else mislead the user perception of ongoing operations.

Aware avoids this problem by displays *Security Messages* to users on the *Status Bar*, a reserved portion of the screen drawable only by the *WindowManager* system service, a privileged process part of the Android OS. A similar approach has been adopted by Bianchi *et al.* [28], where the *Navigation Bar* has been used to host a security indicator as solution against User Interface (UI) deception. However, Bianchi's solution, as it is, cannot be adopted to provide visibility to users about operations targeting I/O devices or to automatically prevent operation programmatically initiated by processes run-

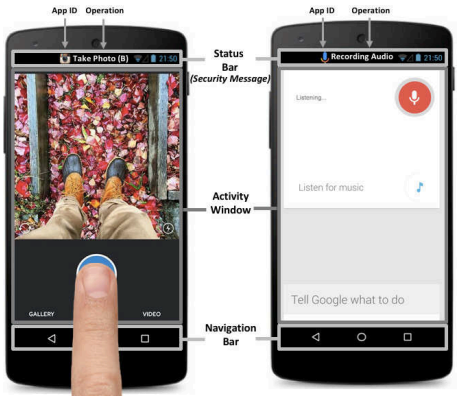

(A) Pending Operation (B) Ongoing Operation

Figure 4: Security Messages displayed on the Status Bar

ning apps, without users interaction. In fact, Bianchi's solution does not provide the necessary mechanisms to bind users interaction to access requests for I/O devices from processes running apps.

Aware uses *Security Messages* displayed on the *Status Bar* to convey, to users, two types of events: (1) *pending operations* initiated by users; and (2) status feedback about *ongoing operations* authorized by users. A *Security Message* includes the app identifier (e.g., app icon or name) and a text message specifying the target operation. The first type of message makes users aware of the operation resulting from the interaction with a soft-button displayed by an app on screen. For example, in Figure 4 (A), if the user presses the button depicting a camera, the *Security Message* specifies that the Instagram app will take a photo using the smartphone front camera[8]. The second type of message informs users about ongoing authorized sessions targeting on-board I/O devices. As example, in Figure 4 (B), a *Security Message* is used to inform the user that the Google Voice Search app is using the microphone to listen to the user voice for commands. If multiple operations are simultaneously targeting different I/O devices, the *Security Message* alternate messages to make users aware of all the ongoing operations.

Security Messages are used to validate precondition P3, ongoing condition O1 and exit condition E3. Satisfying these conditions is sufficient to reliably provide visibility to users over sensitive operations targeting I/O devices, therefore guaranteeing SP3 and SP5.

Aware *Security Messages* are always visible to users even if apps are in full screen mode. Upon use-initiated operations targeting I/O devices (i.e., press soft-button to take a photo), Aware systematically reactivates the *Status Bar* to display a *Security Message* specifying the pending operation. Thus, any attempt by malicious apps to draw a fake *Status Bar* with a fake *Security Message* would fail, since the original *Status Bar* is always drawn on screen.

[8](F) used to indicate front camera and (B) indicate back camera.

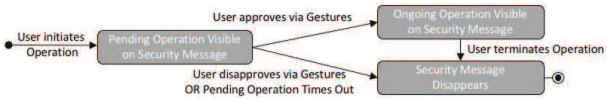
User initiates Operation
Pending Operation Visible on Security Message
User approves via Gestures
Ongoing Operation Visible on Security Message
User terminates Operation
User disapproves via Gestures OR Pending Operation Times Out
Security Message Disappears

Figure 5: State Transition Diagram for Security Messages based on User-Initiated Interactions

6.3 Eliciting User Input for Approval

The requirement for users to approve or abort pending app requests for operations on I/O devices by providing user input through GUI elements, guarantees SP2. On-screen prompts could be used to request approval, every time I/O devices are accessed by a process running an app, in response to user-initiated interactions. However, while prompts attempt to verify users' intention, in practice, they create an excessive burden on users, which leads to users ignoring these prompts eventually[9]. Therefore, prompting users every time I/O devices are accessed seems unreasonable.

To avoid excessive burden on users and, at the same time, enforce a per-access approval, Aware uses a *Gesture Identification* mechanism to identifying specific sequence of gestures, by users, on smartphones' screen. Gestures are intercepted and analyzed, in real-time, to infer users' intention when interacting with apps. Captured gestures on screen are mapped with undergoing operations performed by apps running in foreground. User-initiated interactions are combined with *Security Message* on screen, as depicted in the state machine diagram in Figure 5. With Aware, operations targeting I/O devices can *only* be initiated by user, pressing and holding down a soft-button on screen. Upon user-initiated operations, Aware displays the pending operation on a *Security Message*, for a preset time period, after which the operation is abort in absence of user interaction[10]. After looking at the *Security Message* users can confirm the operation by simply releasing the soft-button, or aborting the pending operation by sliding their finger out from the soft-button area, as shown in Figure 6 (A).

We could have designed Aware in order to have two different areas of the screen where users could place their finger to either deny or allow operations, as shown in figure 6 (B). This solution would have been subject to social-engineering attacks because the two areas would appear in the *Activity Window*, allowing malicious apps to overlay fake messages, and swap the deny area with the allow area to trick the user into allowing an operation.

The Aware *Gesture Identification* mechanism also support an alternative method that makes use of the fingerprint scanner to authenticate users interacting with

[9]On average, there are 8 requests per minute by processes running apps to request permission to access sensitive resources [36].

[10]The timer is used to support apps that require users to keep pressing down a button to perform the operation (i.e., record a video).

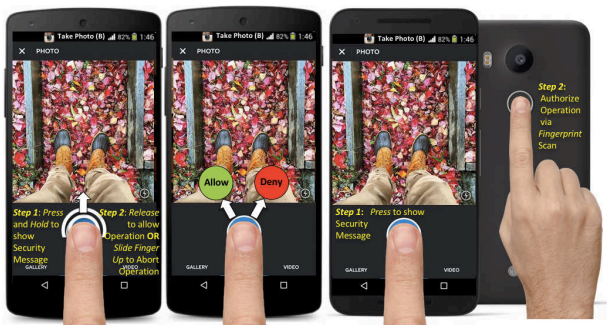
Step 1: Press and Hold to show Security Message
Step 2: Release to allow Operation OR Slide Finger Up to Abort Operation
Allow
Deny
Step 1: Press to show Security Message
Step 2: Authorize Operation via Fingerprint Scan

(A) (B) (C)

Figure 6: Alternative Approaches to retrieve Users Input

smartphonesUsers scan their finger to confirm pending operations displayed on *Security Messages*, as illustrated on the right side of Figure 6 (C). Aware interprets the absence of specific sequences of gestures, from users, how operations not matching users' intention and volition, therefore, blocks and logs attempts from malicious apps trying to programmatically activate security-sensitive operations targeting on-board I/O device. The *Gesture Identification* mechanism is used to validate precondition P4, which in conjunction with P1, P2 and P3 are sufficient to guarantee that operations performed over on-board I/O devices match users' intention and volition, therefore guaranteeing SP2.

For users willing to lower the security of their mobile platforms, Aware allows to disable the *Gesture Identification* mechanism per-app or when a remote controller (i.e., Bluetooth selfie stick) is used. However, we discourage white listing apps, even after a certain period of usage, because apps can dynamically change their behavior during time, due to automatic, periodic, software updates. Furthermore, apps could ask another apps to perform specific operations targeting I/O devices, through the intent mechanism. Thus, a white listed app could be tricked in serving a request coming from a malicious app.

6.4 Supporting Retrospective Actions

The requirement to log actions occurring during the execution of operations targeting I/O devices, guarantees SP4. To support retroactive actions by users, Aware generates three type of access logs for security-sensitive operation targeting on-board I/O devices. First, Aware logs any failed attempt by running processes in accessing I/O devices, due to lack of necessary conditions required to allow requested operations. These logs are accessible in the *Blocked Accesses* section, shown in Figure 7, and allow users to identify apps that attempts to perform stealthy operations while running as background services. Second, Aware logs any operation denied by users though the *Gesture Identification* mechanism. These logs are accessible in the *Denied Accesses* section, shown in Figure 7, and allow users to identify apps using social-engineering techniques to trick them in authorizing undesired operations. Third, Aware log any operation performed over I/O devices, authorized by users though the *Gesture Identification* mechanism, allowing users to track authorized operations.

To better catch users' attention, attempted access violations are signaled by Aware by producing a sound and showing a *Security Message* communicating undesired behaviors from running apps. The Aware *Logs* can be accessed by users anytime, from the app menu or by tapping on the *Security Message* displayed on the *Status Bar*. Each access log entry reports information regarding apps ID, date, time and operations performed by apps, as shown in Figure 7.

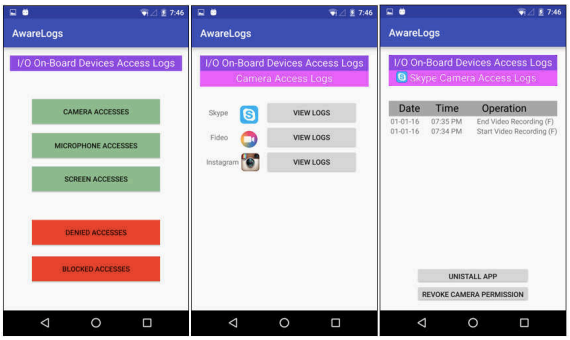

Figure 7: On-Board I/O Devices Access Logs

Aware *Logs* allow users to perform two retrospective actions: (1) uninstall apps identified as malicious, and (2) revoke granted permissions to prevent future undesired accesses[11]. Retrospective actions can be taken by users either immediately, when the I/O devices is still being used by apps to block undesired operations, or after reviewing what apps are doing over time.

The Aware *Log* mechanism is used to validate Ongoing Condition O2 and Exit Condition E2. Satisfying these conditions is sufficient to allow users to perform retrospective actions over sensitive operations targeting I/O devices, therefore guaranteeing SP4.

7 Implementation

We have implemented a prototype of Aware by modifying a recent release[12] of standard Android OS (version android-6.0.1_r5) available through the Android Open Source Project (AOSP) [6]. The Aware prototype includes new components and modifies some pre-existing system services, libraries and the SystemUI app. Their footprint is about 325 LOC in C, 680 LOC in C++ and

[11] This mechanism supports the Android AppOps mechanism reintroduced starting from Android OS 6.0 Marshmallow [4], which allows to revoke permissions granted at install time, to running apps.

[12] A script to automatically integrate Aware on top of previous versions of Android framework components and libraries is also available. Aware source code will be made available on github.com.

882 LOC in Java. We tested the Aware prototype on a Nexus 5 and Nexus 5X[13] smartphones. In the following paragraphs, we briefly describe the new and modified Aware components.

Aware Hooks are placed inside the original *AudioSystem, MediaServer* and *SensorManager* system services to capture the acquisition and release of I/O devices and the reading of sensor data. Other hooks are placed inside the original *InputManager* and *GestureDetector* to capture users' input events. Hooks retrieve the PID of the calling processes, operations requested *Opr* and target devices *Dev*, which are then passed as a parameters in a call back to the Aware Conditional Engine.

Aware Security Messages are implemented by modifying the *SurfaceManager, SurfaceFlinger, WindowManager* and the *SystemUI*. In particular, an ImageView has been added to the Status Bar to display app IDs. Furthermore, a TextView has been added to display ongoing operations, and the service handling the status bar has been modified to receive and handle intent actions sent by the Aware *Conditional Engine*.

The *Aware Conditional Engine* is a new component added to the original Android framework. It includes: (1) the *Callback Handler* in charge of processing callbacks from Aware hooks, (2) the *Verification Service*, in charge of validating preconditions, ongoing conditions, and exit conditions based on the information retrieved from callbacks by Aware *Hook*; and (3) the *Conditional Rule Store*, designed to store and retrieve conditional rules used to enforce control over I/O devices.

The *Aware Gesture Identification* module is implemented by modifying the *InputManager, NativeInputManager, FingerprintManager, InputFlinger* and *GestureDetector*, to translate raw input data into higher-level events (key press, gesture, and such), and then propagates them to the *Conditional Engine* through callbacks.

8 User Study and Aware Evaluation

8.1 Study Objectives

We performed a comprehensive laboratory-based survey, user study, and system experiments with the following five objectives. First, we survey users' privacy and security attitudes as they pertain to the malicious use of I/O devices, and investigate users' awareness about RAT attacks. Second, we observe users' vigilance during a series of interactive tasks, while RAT attacks targeting on-board I/O devices are deployed and the Aware defense mechanisms are not active. Third, we investigate the effectiveness of Aware from two perspectives. Initially, during another series of interactive tasks, we observe whether users notice and can adequately respond to customized social-engineering attacks, when those can be thwarted with the user interface components of Aware. We then debrief users about their experience with Aware. Fourth, we measure whether Aware effectively and systematically shields users from practically deployed RAT attacks, while they perform the second series of interaction tasks. Fifth, we measure the performance overhead Aware incurs on the critical paths of processing app requests.

[13] Equipped with a fingerprint scanner.

8.2 Study Components

The study has three survey components, two interactive user task sequences, and one group of measured tasks. We obtained IRB approval at our institution.

Surveys: Individuals completed an initial questionnaire with demographic questions and questions about their usage of mobile platforms. A second survey debriefed participants about the first series of interactive tasks, performed using an of-the-shelf Android smartphone and investigated their privacy and security attitudes. A third survey debriefed participants about the second series of interaction tasks, performed using a smartphone running Aware, and their perceptions about Aware. The surveys included standard Likert-type psychometric scale questions (e.g., to measure attitudes) as well as open-ended question formats (e.g., to solicit participants' experiences during the interactive tasks). Surveys were implemented on Qualtrics and executed on a lab computer.

Interactive Tasks: In the first series of interactive tasks, we studied participants' potential reactions to practically deployed RAT attacks. Participants were asked to interact with a Nexus 5 smartphone, running a vanilla version of the Android OS (6.0.1_r25), and to perform 9 tasks ranging from taking a picture with the smartphone's camera to sending an email. The first 5 tasks (T1-T5), summarized in Table 3, were not associated with any RAT attacks. Tasks T6-T10 were associated with 4 different, visibly noticeable attacks (A1-A4), also summarized in Table 3. These attacks were carefully triggered by the experimenter, while participants engaged in the interactive tasks. The attacks varied in the degree to which they are perceivable, as highlighted in Table 3. Please note that individuals were not explicitly instructed about the presence of RAT attacks before the tasks, however we asked them to report unusual behaviors verbally and in the survey.

Before the second series of interactive tasks, subjects were debriefed about the previously experienced attacks. We further familiarized them with the Aware system through instructional materials and by allowing them to inspect the Aware user interface on a Nexus 5X smartphone. As before, the participants were engaged in several interactive tasks. First, participants engaged in tasks T1-T5 and T10, which did not present any *noticeable*

Table 3: List of Experimental Tasks and Attacks Description

Task ID	Task Description	App Used	Attack Source	Attack Description	Perceivable	Detection Rate
T1	Take a picture	Instagram (B)	None	N/A	N/A	N/A
T2	Take a video	Fideo (B)	None	N/A	N/A	N/A
T3	Record a voice message	Messenger (B)	None	N/A	N/A	N/A
T4	Record a video message	Skype (B)	None	N/A	N/A	N/A
T5	Record the device screen	Rec. (B)	None	N/A	N/A	N/A
T6-A1	Navigate Internet	Chrome (B)	Krysanec (M)	Stealthy Photo	Camera Shutter Sound	18%
T7-A2	Watch a Video	YouTube (B)	Soundcomber (M)	Stealthy Voice Recording	None	0%
T8-A3	Add a new contact	Contacts (B)	Dendroid (M)	Stealthy Video Recording	UI Slow Down	1%
T9-A4	Send email	Gmail (B)	PlaceRaider (M)	Stealthy Photos	UI Slow Down	0%
T10	Take a screenshot	None*	None	N/A	N/A	N/A
T11-A5	Record a video	SimpleFilters †	SimpleFilters †	Stealthy Screenshot	Security Message Mismatch	82%
T12-A6	Take a photo	SimpleFilters †	SimpleFilters †	Stealthy Voice Record	Security Message Mismatch	76%
T13	Analyze Aware Logs	Aware Logs (B)	None	N/A	N/A	N/A

(B) Benign App (M) Malicious App * Performed by pressing the power and volume-down physical buttons at the same time
† SimpleFilters appears as a benign app, but includes functionality to run additional I/O operations beyond those consented by users

RAT behaviors. Then, in tasks T11 and T12, we used a test RAT app to investigate users' responses to attacks where the user consented to one action, but the app performed additional, unapproved actions. These social-engineering attacks (A5 and A6) aimed to trick users into executing unwanted I/O operations, such as recording voice when they consented to the app taking a photo. Using the Aware security message and gesture mechanism, we investigated whether participants could notice and thwart the attacks. The series of tasks concluded with T13, which did not include any attacks. We did not brief individuals about which tasks were associated with attacks or which I/O devices would be targeted. We summarize the 9 tasks and associated attacks in Table 3.

We recorded participants' interactive behaviors, and their survey responses. In addition, we applied a think-aloud protocol by encouraging participants to speak out about their experiences while being engaged in the tasks.

Measured Evaluation: During the second series of tasks, participants were still exposed to attacks programmatically and persistently triggered from the four RAT apps used in our study (i.e., Krysanec, Soundcomber, Dendroid, and PlaceRaider). The Aware system was expected to shield the user from these attacks, and therefore no *noticeable* effects should have been observable by the participants. As such, the purpose of the measurement task is to evaluate whether Aware effectively and automatically shields the participants *while they engage in realistic user interaction behaviors*.

All participants completed the entire set of study components in about 25-35 minutes and were compensated with a $10 gift card.

8.3 Results

Demographics and Mobile Platforms Usage: In total, 74 participants completed the whole set of surveys and tasks. The majority of the sample were between 20-29 years old (76%). We recruited predominantly undergraduate and graduate students; the majority having an international background (70%), and fields of study *different* from computer science (75%). Most participants actively used smartphones (99%) and additional devices associated with third-party apps such as tablets (54%), and to a lesser extent smart watches and fitness bands (10%).

Privacy Attitudes: We asked participants how concerned they are about threats to their personal privacy when using smartphones, and found that 43% were moderately or extremely concerned. Participants were even more concerned about privacy and security aspects as they related to third-party apps (57%). Most important to our study, concern levels were high for the misuse of smartphones' camera (62%) and microphone (55%).

RAT Awareness and Security Behaviors: The majority of the participants stated that they were aware that apps could access the camera (56%) and the microphone (56%) of their smartphones at any time without repeatedly asking for consent. However, participants had little knowledge of specific RATs that exploit smartphones' I/O devices, such as Dendroid and SoundComber (4% each), and Krysanec and PlaceRaider (3% each). A small number of participants (8%) were able to articulate how malicious apps apply social-engineering techniques to misled users into taking an action. Only 24% of the participants use a mobile anti-malware product, whereas 78% of subjects stated that they avoid downloading apps from unofficial app stores.

Identification of Threats without Aware: The attacks in the first series of interaction tasks varied in the degree to which *risk signals* in the vanilla Android system are perceivable by a user. A1 was associated with the camera shutter sound when the Krysanec malware took a stealthy photo, while participants were browsing the Internet. Only 18% correctly noticed that a camera shutter sound was audible. 8% incorrectly thought that a screenshot was taken. 4% merely noticed a sound. Not a single participant stated any suspicion in the survey or the think-aloud comments that malware or a security problem could be responsible for the sound. Two participants questioned whether "something" or "someone was taking a screenshot". Only one participant noticed the UI

slow-down (i.e., a perceived glitch in the UI) in task T8-A3 when a stealthy voice recording was initiated by the SoundComber app. None of the participants identified any suspicious activity when being involved in attack tasks T7-A2 and T9-A4. An overview of the detection rates for the RAT attacks is available in Table 3.

In summary, audible risk signals (associated with the built-in camera in A1) are noticed by at least some participants, but their interpretation of the risk, ability to identify the cause or at least to question the cause of the unexpected behavior is highly limited. State-of-the-art malware apps engaging in stealthy voice and video recordings remained almost entirely unnoticed. PlaceRaider (A4) which also takes stealthy pictures and sidesteps audible risk signals was not noticed by any participant.

Interaction of Users with Aware: Using Aware substantially simplifies the problem space of unwanted I/O access in two ways. First, any access to an I/O device which is not associated with a user interaction will be automatically blocked. Thereby, attacks that participants were encountering in the first series of interactive tasks would be automatically thwarted. We will verify this aspect of Aware's effectiveness in the subsequent measurement results subsection. Second, only user-intended interactions which explicitly trigger an I/O device will be allowed to proceed. Note, however, that malicious apps could nevertheless try to hijack these user-intended activities. Aware also aims to thwart attacks that coincide with such user-intended I/O interactions. For this purpose, Aware provides user interface features (i.e., Aware security messages and gesture identification) that help to *notice* and to *block* such attacks. In addition, users can *inspect the Aware log* to verify which wanted or unwanted I/O operations have taken place.

To investigate the effectiveness of Aware when users encounter such sophisticated attacks (which exploit the *context* of users' interactions), we have customized a simple photographic filter app called SimpleFilters. The app is modified to attempt to take a screenshot (A5) and a voice recording (A6) stealthily when users intentionally engage in a different type of I/O interaction. Using Aware, 82% of the participants noticed the attack while performing task T11-A5, and 76% noticed the attack while performing task T12-A6. In contrast to the first series of user interactions, these participants were also able to clearly articulate what attacks took place. A typical debriefing message for attack A5 was: "It was taking a screenshot, while I asked it to take a video", whereas for attack A6, messages were a variations of: "Aware message said the microphone was being accessed".

Inspecting the system logs, we were also able to determine how often participants used the Aware gestures to *block* the attacks that they noticed. For A5, all of the 82% of the participants who noticed the attack successfully used the gestures to abort the task. Similarly, all participants who noticed attack A6 succeeded in blocking the attempt to record audio instead of taking a picture. In the final task (T13), we asked individuals to inspect Aware Logs to evaluate which I/O access operations had taken place during the second series of interactions. 88% of all participants found Aware Logs helpful in identifying suspicious activities from running apps, and they were clearly able to articulate what attacks had taken place.

After the interaction tasks, we solicited further feedback from the participants. 90% of the participants found Aware more secure than the vanilla Android OS, and 80% found it as (or more) usable compared to the vanilla Android OS. These are encouraging results since additional security mechanisms often meet with user resistance, for example, because they may distract from the user's primary task. Further, 57% of the participants said they would prefer the Aware notice and gesture mechanism compared to other notification options. For example, only 21% of the participants preferred to be prompted with a permission dialog at every access. Further, only 10% of them stated that they would prefer to be asked for permission at install time, and 8% of them at first use. Most importantly, 99% of the participants would like Aware integrated in their current mobile OS.

Measurements to Evaluate Aware: To evaluate the effectiveness of Aware in the presence of realistic user interactions, we allowed the 4 RAT apps used in the first series of user interaction tasks to also be active during the second series of interaction tasks. We also tested whether activities from the customized SimpleFilters app would be blocked if they did not coincide with the users' interactions with the I/O devices (i.e., in T11 and T12). In order to monitor whether any of the malicious activities of these RAT apps were successful, we used logcat [4], the Android logging system, which provides a mechanism for collecting system debug output about activities from various apps and system services. For the 74 sessions involving participants, we found that RAT apps attempted 1080 times to perform stealthy operations targeting I/O devices, but they *never* succeeded in accessing the onboard camera, microphone or screen content, as result of systematically validating preconditions P1 and P4. In other words, the absence of users' interaction and consent prevented RAT apps from succeeding in performing stealthy operations while running services in the background. Furthermore, based on the logs, there were no run-time exceptions, triggered by the Aware components, which could have caused any of the 9 well-known legitimate apps to crash or unexpectedly terminate.

Summary[14]*:* Aware prevented all attempts from RAT apps to perform stealthy operations that did not coin-

[14] See Appendix D for a summary of selected results.

Table 4: Aware Performance Overhead in μs. Numbers give the mean value and corresponding standard deviation after 10,000 runs.

	Vanilla Android OS		Aware		
	Nexus 5	Nexus 5X	Nexus 5	Nexus 5X	Max (Min) Overhead
Front Camera	15.90±1.54	14.39±1.12	16.11±1.77	15.01±1.38	4.04% (2.21%)
Back Camera	16.08±1.32	15.68±1.87	16.44±1.06	16.37±1.91	4.31% (2.57%)
Microphone	12.36±2.01	11.86±1.99	12.65±2.15	12.32±1.85	4.03% (2.19%)
Screen	17.76±0.99	16.23±0.69	18.61±0.90	17.02±1.01	4.79% (2.94%)

cide with users' intended I/O access operations and considerably reduced the success rate of social-engineering attacks, without breaking any apps' logic. Therefore, Aware significantly raises the bar compared to the detection rate of state-of-the-art static/dynamic analysis tools, and anti-malware tools, available to users to identify malicious apps running on smartphones[15]. We anticipate that with additional experience users will become even more proficient with the Aware security messages and the gesture mechanism, which would further reduce the effectiveness of social-engineering techniques. However, the achieved results are very impressive given the sophisticated nature of the attacks tested with the SimpleFilters app[16]. Aware automatically blocks all attacks which are not carefully socially-engineered and significantly reduces the attention burden placed on users, thereby, reduces habituation and notice fatigue [38].

8.4 Performance Evaluation

We have measured the overhead introduced by Aware while handling each access request for operating on-board I/O devices, such as the camera to take photos and video, the microphone to record audio, and the screen to capture screenshots. Due to lack of publicly available benchmarks for Android OS, we only provide microbenchmark analysis of such operations for two phones, a Nexus 5 and a Nexus 5X running Android OS (version android-6.0.1_r5). The overhead is calculated by measuring the time interval from the time the request is made by the process running the app to the time the operation is granted/denied by Aware. Table 4 reports the average time over 10,000 requests, the standard deviation and the maximum recorded overhead introduced by Aware. Overall, Aware introduces a negligible overhead of the order of 1 μs per access. The maximum recorded overhead is 4.79% while accessing the screen buffers.

9 Related Work

User-Driven Access Control (UDAC) [25] attempts to include users in the access control decision loop. Application level access control gadgets (ACGs) are used to verify that actions requested by apps genuinely comes from users. However, UDAC is subject to social-engineering attacks [28, 27, 16], such as *draw on top*[17] and *app switch*[18], because ACGs are displayed in a portion of the screen accessible to apps.

SemaDroid [34] is a privacy-aware sensor management framework for smartphones that allows users to monitor sensors usage by installed apps, and control the disclosure of sensed data. However, users are supported in identifying undesired accesses only after apps have already performed operations targeting I/O device.

Petracca *et al.* [35] propose AuDroid, an extension to the SELinux reference monitor integrated into the Android OS, to enforce lattice policies over the dynamically changing use of system audio resources (e.g. microphone and speaker). However, AuDroid only deals with accesses to microphone and speaker from third-party apps and does not provide a solid and general solution for other on-board I/O devices and sensors.

Bianchi *et al.* [28] address the problem of malicious apps surreptitiously replacing or mimicking the GUI of other apps to mount phishing and click-jacking attacks. However, Bianchi's solution cannot systematically prevent operations programmatically initiated by processes running apps, without users' interaction, because it does not provide the necessary mechanisms to bind users' interaction to access requests for I/O devices from processes running apps. Moreover, Bianchi's use visibility to transfer more responsibility to users when attack scenarios have to be identified, whereas visibility should be used to support users in making decision only when it is not possible to prevent attacks systematically.

10 Conclusion

In this paper, we presented Aware, a security framework for authorizing app requests to perform sensitive operations using I/O devices, which binds app requests with user intentions to make all uses of certain I/O devices explicit. We evaluated the proposed defense mechanisms through laboratory-based experimentation and a user study, involving 74 human subjects, whose ability to identify undesired operations targeting I/O devices increased significantly. Without Aware, only 18% of the participants were able to identify attacks from tested RAT apps. Aware systematically blocked all the attacks, in absence of user-initiated interaction, and supported users in identifying 82% of more sophisticated attacks, which used social-engineering techniques to hijack user-initiated operations. Aware introduced only 4.79% maximum performance overhead over operations targeting I/O devices.

[15] Detailed analysis results reported in Appendices B and C.

[16] Research on carefully crafted Phishing attacks shows that, even with repeated security training, a significant share of users will fall for such attacks [37].

[17] Drawing of graphical elements over other apps.

[18] Malicious app replaces the legitimate top Activity with one of its own.

References

[1] Dendroid malware can take over your camera, record audio, and sneak into Google Play. https://blog.lookout.com/blog/2014/03/06/dendroid/

[2] Smartphone OS Market Share. http://www.idc.com/prodserv/smartphone-os-market-share.jsp

[3] Krysanec Trojan: Android backdoor lurking inside legitimate apps. http://www.welivesecurity.com/ 2014/08/12/krysanec-trojan-android

[4] Official Android Documentation. http://developer.android.com

[5] SEAndroid. http://seandroid.bitbucket.org

[6] Android Open Source Project. https://source.android.com/

[7] Java Native Interface. http://en.wikipedia.org/wiki/Java Native Interface

[8] Google Play app store. https://play.google.com/

[9] Android Central. http://www.androidcentral.com/stagefright

[10] Apks Free app store. http://www.androidapksfree.com

[11] Mobile Apk World app store. http://mobileapkworld.com

[12] VirusTotal. Free Online Virus, Malware and URL Scanner. https://www.virustotal.com

[13] Smartphones Are Used To Stalk, Control Domestic Abuse Victims. http://www.npr.org/sections/alltechconsidered/2014/09/15/346149979/smartphones-are-used-to-stalk-control-domestic-abuse-victims

[14] Google Bouncer. http://blog.trendmicro.com/trendlabs-security-intelligence/a-look-at-google-bouncer/

[15] Speed up your car insurance claim. https://www.esurance.com/photo-claims

[16] Social-Engineering Attacks. https://en.wikipedia.org/wiki/Social_engineering_(security)

[17] PNC Mobile Banking. https://www.pnc.com/en/personal-banking/banking/online-and-mobile-banking/mobile-banking.html

[18] Samsung Know White Papers. https://www.samsungknox.com/en/support/knox-workspace/white-papers

[19] Yee, Ka-Ping. Aligning security and usability.

[20] Park, Jaehong and Sandhu, Ravi. The UCON ABC usage control model.

[21] Anderson, James P. Computer Security Technology Planning Study. Volume 2.

[22] Robert Templeman and Zahid Rahman and David Crandall and Apu Kapadia. PlaceRaider: Virtual theft in physical spaces with smartphones.

[23] Schlegel, Roman and Zhang, Kehuan and Zhou, Xiao-yong and Intwala, Mehool and Kapadia, Apu and Wang, XiaoFeng. Soundcomber: A Stealthy and Context-Aware Sound Trojan for Smartphones.

[24] Felt, Adrienne Porter and Ha, Elizabeth and Egelman, Serge and Haney, Ariel and Chin, Erika and Wagner, David. Android permissions: User attention, comprehension, and behavior.

[25] Roesner, Franziska and Kohno, Tohru and Moshchuk, Alexander and Parno, Bryan and Wang, Harry Jiannan and Cowan, Crispin. User-driven access control: Rethinking permission granting in modern operating systems.

[26] World-driven access control for continuous sensing Roesner, Franziska and Molnar, David and Moshchuk, Alexander and Kohno, Tadayoshi and Wang, Helen.

[27] Anderson, Fraser and Grossman, Tovi and Wigdor, Daniel and Fitzmaurice, George. Supporting Subtlety with Deceptive Devices and Illusory Interactions.

[28] Bianchi, Antonio and Corbetta, Jacopo and Invernizzi, Luca and Fratantonio, Yanick and Kruegel, Christopher and Vigna, Giovanni. What the App is That? Deception and Countermeasures in the Android User Interface.

[29] Smalley, Stephen and Craig, Robert. Security Enhanced Android: Bringing Flexible MAC to Android.

[30] Nadkarni, Adwait and Enck, William. ASM: A programmable interface for extending Android security.

[31] Backes, Michael and Bugiel, Sven and Gerling, Sebastian and von Styp-Rekowsky, Philipp. Android Security Framework: Extensible multi-layered access control on Android.

[32] Chen, Kai and Wang, Peng and Lee, Yeonjoon and Wang, XiaoFeng and Zhang, Nan and Huang, Heqing and Zou, Wei and Liu, Peng. Finding unknown malice in 10 seconds: Mass vetting for new threats at the Google-Play scale.

[33] Tam, Kimberly and Khan, Salahuddin and Fattori, Aristide and Cavallaro, Lorenzo. CopperDroid: Automatic Reconstruction of Android Malware Behaviors.

[34] Xu, Zhi and Zhu, Sencun. SemaDroid: A Privacy-Aware Sensor Management Framework for Smartphones.

[35] Petracca, Giuseppe and Sun, Yuqiong and Atamli, Ahmad and Jaeger, Trent. AuDroid: Preventing Attacks on Audio Channels in Mobile Devices.

[36] Primal Wijesekera and Arjun Baokar and Ashkan Hosseini and Serge Egelman and David Wagner and Konstantin Beznosov. Android Permissions Remystified: A Field Study on Contextual Integrity.

[37] Sheng, Steve and Holbrook, Mandy and Kumaraguru, Ponnurangam and Cranor, Lorrie Faith and Downs, Julie. Who Falls for Phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions.

[38] Böhme, Rainer and Grossklags, Jens. The Security Cost of Cheap User Interaction.

Appendices

A Android Permission Set Analysis

The analysis of 74 apps from third-party app stores [11, 10] and 329 apps from the official Google Play [8], shows concerning results, summarized in Table 5. In particular, many of the analyzed apps could potentially behave as RAT, since they have the necessary permissions to perform stealthy operations targeting on-board I/O devices. For example, from the Google Play, 83.89% of apps could potentially take stealthy screenshots. Furthermore, 25.68% of apps from third-party app stores could potentially take stealthy photos. In each cell of Table 5, the first value represents the number, the second value the percentage of apps, among all the app analyzed in the same category, that have the permissions required to perform the stealthy operation specified in the first column. We are not aware if these apps are actually misusing their permission, but we want to point out that it is possible for these apps to misuse their permissions to perform stealthy operations, and by statically analyzing the set of Android permissions used by apps, it is by no mean sufficient to distinguish between purely benign apps and malicious apps.

B Anti-Malware Tools Detection Analysis

The analysis results relative to RAT detection by the 15 most popular anti-malware apps, available on Google Play [8] and used on smartphones by millions of user

Table 5: Market Apps that potentially could behave as RATs and perform stealthy operations

Potential Feature\Category	Game	Business	Book	Comics	Communication	Education	Entertainement	Total
Third-Party App Stores								
Stealthy Screenshots	18 (90.00%)	11 (68.75%)	1 (100.00%)	2 (100.00%)	20 (95.24%)	0 (0.00%)	9 (69.23%)	61 (82.43%)
Stealthy Photos	0 (0.00%)	4 (25.00%)	0 (0.00%)	1 (50.00%)	13 (61.90%)	0 (0.00%)	1 (7.69%)	19 (25.68%)
Stealthy Videos	0 (0.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)	12 (57.14%)	0 (0.00%)	0 (0.00%)	12 (16.22%)
Stealthy Audio	2 (10.00%)	0 (0.00%)	0 (0.00%)	0 (0.00%)	12 (57.14%)	0 (0.00%)	4 (30.77%)	18 (24.32%)
Number of Apps	20	16	1	2	21	1	13	74
Google Play Store								
Stealthy Screenshots	38 (95.00%)	33 (91.67%)	25 (83.33%)	30 (90.91%)	61 (87.14%)	52 (67.53%)	37 (86.05%)	276 (83.89%)
Stealthy Photos	3 (7.50%)	9 (25.00%)	2 (6.67%)	4 (12.12%)	32 (45.71%)	4 (5.12%)	7 (16.28%)	61 (18.54%)
Stealthy Videos	0 (0.00%)	2 (5.56%)	1 (3.33%)	1 (3.03%)	27 (38.57%)	1 (1.30%)	3 (6.98%)	35 (10.64%)
Stealthy Audio	0 (0.00%)	3 (8.33%)	2 (6.67%)	3 (9.09%)	30 (42.96%)	6 (7.79%)	4 (9.30%)	48 (14.59%)
Number of Apps	40	36	30	33	70	77	43	329

around the world, are summarized in Table 6. The *Installs* column indicates the number of installs performed by users on their smartphones. The *Reviews* column specifies how many people gave a personal review and a score for the anti-malware app on Google Play. Finally, the *Score* column reports the average score received by the app over a scale of 5 by user reviewing the app itself.

All the anti-malware apps have been updated with the most recent malware database before starting the scan, indeed 3 anti-malware apps (marked with $\star$ in Table 5) have detected the Stagefright vulnerability [9] that would allow malicious code to send fraudulent MMS, only recently discovered. The analysis results are based on a first scan before malware installation, and a second scan during the execution of the malware, when the anti-malware has been kept actively scanning for 10 minutes. Subsequently, three consecutive scans have been performed, after malware installation. After the first scan, the anti-malware has been configured to actively keep scanning. We made sure to select full/deep scan from the scanning options. The three successive scans have been manually activated to force the anti-malware to rescan the entire system. At each new malware installation the smartphone has been flashed again with a clean copy of the OS and anti-malware software installation.

The analysis revealed that most anti-malware tools are able to detect well-known RATs (e.g., Dendroid and Krysanec). We believe that this is due to the fact that well-known RATs have been classified and a signature has been generated and distributed on the Web. On the other hand, proof-of-concept RATs (e.g., PlaceRaider, SoundComber and StealthyStalker) are unknown and gone undetected by anti-malware tools even though they use similar techniques used by well-known RATs. Exceptionally, the AVAST Mobile Security Anti-Malware identifies some malice in both PlaceRaider and StealthyStalker.

On the Android OS side, an interesting finding was that, at install time, an alert was triggered to block the installation of Krysanec. At the second attempt, the Android OS asked the user if to proceed with the installation anyhow. Additionally, while uninstalling the app, Krysanec attacked the operating system by exploiting privilege escalation.

C Static and Dynamic Analysis Tools

The analysis results relative to RAT detection by four state-of-the-art static and dynamic analysis tools are reported in Table 6.

VirusTotal [12], originally developed by Hispasec and now own by Google, is a free service that analyzes suspicious files and URLs and facilitates the quick detection of viruses, worms, trojans, and all kinds of malware. It uses 56 different anti-malware products and 61 online scan engines to check for viruses. VirusTotal was selected by PC World as one of the best 100 products of 2007. As shown in Table 6, VirusTotal detects well-know RATs (e.g., Dendroid and Krysanec). The two RATs are identified as malicious with a score of respectively 22/56 and 20/56. The nominator in the score fractions refers to the number of tools that identify the app as potentially malicious, the denominator indicates the total number of tools that have analyzed the app.

MassVet [32] compares a submitted app with apps already on a market, focusing on the difference between those sharing a similar UI structure (indicating a possible repackaging relation), and the commonality among those seemingly unrelated. MassVet uses a "DiffCom" analysis on top of an efficient similarity comparison algorithm, which maps features of an app's UI structure or a method's control-flow graph to a value for a fast comparison. As shown in Table 6, MassVet detects malicious code in 3 of the 5 RATs analyzed.

Google Bouncer [14] is a codename used by Google, for a security service introduced early in 2012, to keep malicious apps off the official Google Play[19]. Bouncer

[19] According to Google, Bouncer was responsible for a 40% drop in the number of malicious apps in its app store.

Table 6: RAT Detection by the 15 Most Popular Android Anti-Malware Tools

Legend: ✓ Malware Detected ✗ Malware Undetected ○ Malware Detected as Privacy Violation ⊗ Anti-malware crashed during Scan	PlaceRaider	Soundcomber	StealthyStalker	Dendroid	Krysanec	Installs	Reviews	Score (Scale of 5)
360 Security Antivirus Boost ⋆	✗	✗	✗	✓	✗	100M-500M	8,092,733	4.6
AndroHelm AntiVirus	✗	✗	✗	✗	✗	100K-500K	5,383	4.1
TrustGo Antivirus & Mobile Security	✗	✗	✗	✓	✓	10M-50M	283,332	4.5
AVAST Mobile Security & Antivirus	✓	✗	✓	✓	✓	100M-500M	3,481,194	4.5
AVG AntiVirus Security Scan ⋆	✗	✗	✗	✓	✓	100M-500M	4,079,893	4.4
Bitdefender Mobile Security & Antivirus	✗	✗	✗	✓	✓	1M-5M	33,909	4.3
Cheetah Mobile Security Antivirus & AppLock ⋆	✗	✗	✗	○	✓	100M-500M	13,258,188	4.7
Dr Web Anti-virus	✗	✗	✗	✓	✓	50M-100M	863,913	4.5
ESET Mobile Security & Antivirus	✗	✗	✗	✓	✓	5M-10M	300,470	4.6
Kaspersky Internet Security	✗	✗	✗	✓	✓	10M-50M	1,019,526	4.6
Lookout Security & Antivirus	✗	✗	✗	✓	✓	100M-500M	830,941	4.4
Malwarebytes Anti-Malware	✗	✗	✗	⊗	✓	1M-5M	53,990	4.3
McAfee Security & Antivirus	✗	✗	✗	✓	✗	10M-50M	334,162	4.3
Symantec Norton Security & Antivirus	✗	✗	✗	✓	✓	10M-50M	509,521	4.4
NQ Security Lab Antivirus Free-Mobile Security	✗	✗	✗	✗	✗	10M-50M	497,640	4.3

Table 7: RAT Detection via Static and Dynamic Analysis

	Static Analysis		Dynamic Analysis	
	VirusTotal	MassVet	Google Bouncer	CopperDroid
PlaceRaider	✗	✗	N/A	N/A
SoundComber	✗	✓	N/A	N/A
StealthyStalker	✗	✓	✗	N/A
Dendroid	✓	✗	N/A	N/A
Krysanec	✓	✓	N/A	N/A

Legend: ✓ Malware Detected ✗ Malware Undetected
N/A Data Not Available

quietly and automatically scans apps (both new and previously uploaded ones) and developer accounts in Google Play with its reputation engine and cloud infrastructure. To test the effectiveness of Bouncer, in detecting RAT apps, we have implemented a proof-of-concept testing app, called *StealthyStalker*[20], able to take stealthy photos, videos, screenshots, record audio and hijack user-initiated operations. To release an app through Google Play, a third-party developer has to participate in Android developer program and submit to Google for review. The app is signed and published by Google only after it passes the review process. As shown in Table 6 and Figure 8, the StealthyStalker app (submitted for publication with the fake name of SimpleFilter) successfully passed the Google Play (Bouncer) review and published, after a couple of hours, despite the hidden malicious code. This means that Google Bouncer did not find any potential harm in the app. Following the ethical hacking practice, we immediately removed the app from Google Play before any user could actually download it, as proved by the download statistic provided by Google. Results for the other 4 RATs are not available due to the fact that we are not authorized to submit code written by other researchers or malicious developers.

[20]Submitted to Google Play under the name of *SimpleFilters*

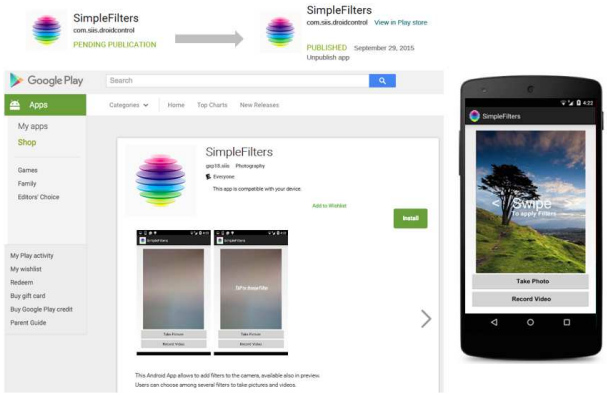

Figure 8: StealthyStalker RAT app published on Google Play under the name of SimpleFilters

CopperDroid [33] is a tool to perform dynamic analysis over Android apps to characterize the behavior of Android malware. It automatically analyzes low-level OS-specific and high-lelvel Android-specific behaviors of Android malware by observing and analyzing system call invocations, including IPC and RPC interactions, carried out as system calls. Although CopperDroid is a powerful tool that dynamically analyzes Android apps, analysis results do not provide any hint on the maliciousness of a given sample. Indeed, by manually analyzing the logs files generated by CopperDroid (e.g, syscalls, pcap, logcat and basicbinder), we were not able to identify evidence of malice for the analyzed software.

D Summary of selected Results from our User Study and Aware Evaluation

Demographic		
Total Participants		74
Age	18-19	6%
	20-29	76%
	30-39	18%
International Students		70%
Nationalities		25
Non Computer Science Major's		75%
Major/Minor's Degrees		34

Mobile Platform Usage	
Smartphones	99%
Tablets	54%
Wearables	10%
Android OS	57%
Apple iOS	42%
Other OS	1%

Participants Concerned about Personal Privacy	
When using Mobile Platfotm	43%
When using third-party Apps	57%
When using Platform's Camera	62%
When using Platform's Microphone	55%

Security Behavior	
Avoid downloading apps from unofficial App Stores	78%
Use a mobile anti-malware product	24%

Participants' Awareness about RAT Operations	
Stealthy accesses to Platform's Camera	56%
Stealthy accesses to Platform's Microphone	56%
Knew about Dendroid or SoundComber	4%
Knew about Krysanec or PlaceRaider	3%

Threats Identification without Aware	
Participants identifying Attack A1	36%
Participants identifying Attack A2	0%
Participants identifying Attack A3	1%
Participants identifying Attack A4	0%

Threats Identification with Aware	
Participants identifying Attack A1	82%
Participants identifying Attack A2	76%

Aware Effectiveness and Usability	
Stealthy Accesses Blocked	100%
User-Initiated Operations Hijacks Prevented	82%
Rated Aware more secure than Original OS	90%
Rated Aware as usable as the Original OS	80%
Prefer Aware Gesture to Access Prompts on Screen	57%

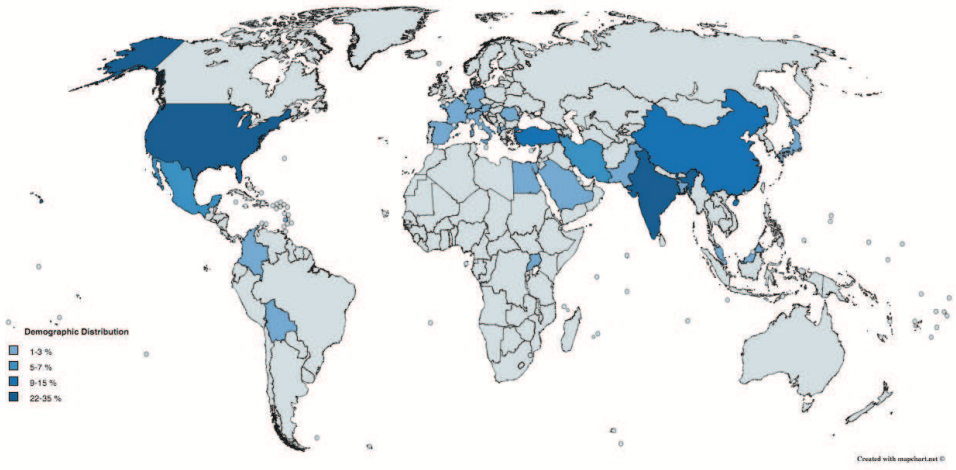

Demographic Distribution of Human Subjects participating in the Study

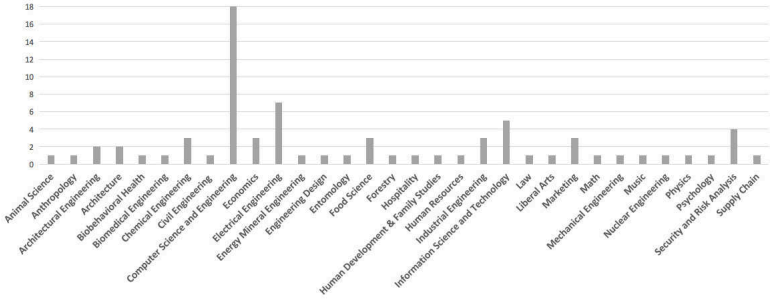

Major/Minor's Degree Distribution of Human Subjects participating in the Study